arXiv:1407.4235v1 [cs.DS] 16 Jul 2014

The List Coloring Reconfiguration Problem for Bounded Pathwidth Graphs

Tatsuhiko Hatanaka, Takehiro Ito, and Xiao Zhou

Graduate School of Information Sciences, Tohoku University,
Aoba-yama 6-6-05, Sendai 980-8579, Japan.
{hatanaka, takehiro, zhou}@ecei.tohoku.ac.jp

Abstract. We study the problem of transforming one list (vertex) coloring of a graph into another list coloring by changing only one vertex color assignment at a time, while at all times maintaining a list coloring, given a list of allowed colors for each vertex. This problem is known to be PSPACE-complete for bipartite planar graphs. In this paper, we first show that the problem remains PSPACE-complete even for bipartite series-parallel graphs, which form a proper subclass of bipartite planar graphs. We note that our reduction indeed shows the PSPACE-completeness for graphs with pathwidth two, and it can be extended for threshold graphs. In contrast, we give a polynomial-time algorithm to solve the problem for graphs with pathwidth one. Thus, this paper gives precise analyses of the problem with respect to pathwidth.

1 Introduction

Graph coloring is one of the most fundamental research topics in the field of theoretical computer science. Let $C = \{1, 2, \ldots, k\}$ be the set of k colors. A (proper) *k-coloring* of a graph $G = (V, E)$ is a mapping $f : V \to C$ such that $f(v) \neq f(w)$ for every edge $vw \in E$. In *list coloring*, each vertex $v \in V$ has a set $L(v) \subseteq C$ of colors, called the *list* of v. Then, a k-coloring f of G is called a *k-list coloring* of G if $f(v) \in L(v)$ holds for every vertex $v \in V$. Figure 1(b) illustrates four k-list colorings of the same graph G with the same list L depicted in Fig. 1(a); the color assigned to each vertex is attached to the vertex. Clearly, a k-coloring of G is a k-list coloring of G for which $L(v) = C$ holds for every vertex v of G, and hence k-list coloring is a generalization of k-coloring.

Graph coloring has several practical applications, such as in scheduling, frequency assignments. For example, in the frequency assignment problem, each vertex corresponds to a base station and each edge represents the physical proximity and hence the two corresponding base stations have the high potential of interference. Each color represents a channel of a particular frequency, and we wish to find an assignment of channels to the base stations without any interference. Furthermore, in list coloring, each base station can have a list of channels that can be assigned to it.

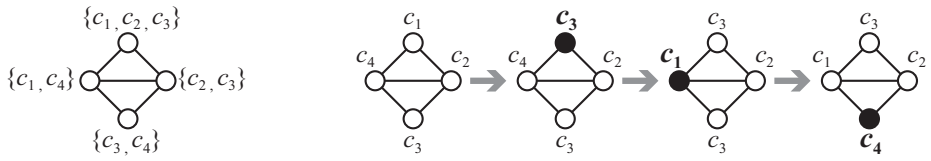

Fig. 1. (a) Graph G and its list L, and (b) a sequence of k-list colorings of G.

1.1 Our problem

However, a practical issue in channel assignments requires that the formulation should be considered in more dynamic situations. One can imagine a variety of practical scenarios where a k-list coloring (e.g., representing a feasible channel assignment) needs to be transformed (to use a newly found better solution or to satisfy new side constraints) by individual color changes (keeping the network functionality and preventing the need for any coordination) while maintaining feasibility (so that the users receive service even during the reassignment).

In this paper, we thus study the following problem: Suppose that we are given two k-list colorings of a graph G (e.g., the leftmost and rightmost ones in Fig. 1(b)), and we are asked whether we can transform one into the other via k-list colorings of G such that each differs from the previous one in only one vertex color assignment. We call this decision problem the k-LIST COLORING RECONFIGURATION problem. For the particular instance of Fig. 1(b), the answer is "yes," as illustrated in Fig. 1(b), where the vertex whose color assignment was changed from the previous one is depicted by a black circle.

1.2 Known and related results

Recently, similar settings of problems have been extensively studied in the framework of reconfiguration problems [13], which arise when we wish to find a step-by-step transformation between two feasible solutions of a problem such that all intermediate solutions are also feasible. This reconfiguration framework has been applied to several well-studied combinatorial problems, including satisfiability [11, 21], independent set [6, 12, 13, 20, 22], set cover, matching [13], shortest path [3, 4, 19], list edge-coloring [14, 16], list $L(2,1)$-labeling [15], and so on.

In particular, the k-COLORING RECONFIGURATION problem (i.e., k-LIST COLORING RECONFIGURATION in which $L(v) = C$ holds for every vertex v) is one of the most well-studied reconfiguration problems, from the various viewpoints [1, 2, 5, 7, 10, 18, 25], as follows.

Bonsma and Cereceda [5] proved that k-COLORING RECONFIGURATION is PSPACE-complete for $k \geq 4$; they also proved that k-LIST COLORING RECONFIGURATION is PSPACE-complete, even for bipartite planar graphs and $k = 4$. On the other hand, Cereceda et al. [10] proved that k-COLORING RECONFIGURATION is solvable for any graph in polynomial time for the case where $1 \leq k \leq 3$.

Then, some sufficient conditions have been proposed so that any pair of k-colorings of a graph has a desired transformation. Cereceda [9] gave a sufficient

condition with respect to the number k of colors: if k is at least the treewidth of a graph G plus two, then there is a desired transformation between any pair of k-colorings of G; the length of the transformation (i.e., the number of recoloring steps) is estimated by Bonamy and Bousquet [2]. Bonamy et al. [1] gave a sufficient condition with respect to graph structures: for example, chordal graphs and chordal bipartite graphs satisfy their sufficient condition.

Recently, Bonsma et al. [7] and Johnson et al. [18] independently developed a fixed-parameter algorithm to solve k-COLORING RECONFIGURATION when parameterized by $k + \ell$, where k is the number of colors and ℓ is the number of recoloring steps. In contrast, if the problem is parameterized only by ℓ, then it is W[1]-hard [7] and does not admit a polynomial kernelization unless the polynomial hierarchy collapses [18].

In this way, even for the non-list version, only a few results are known from the viewpoint of polynomial-time solvability. Furthermore, as far as we known, no algorithmic result has been obtained for the list version.

1.3 Our contribution

In this paper, we study the k-LIST COLORING RECONFIGURATION problem from the viewpoint of graph classes, especially pathwidth of graphs. (The definition of pathwidth will be given in Section 2.)

We first show that the problem remains PSPACE-complete even for graphs with pathwidth two. In contrast, we give a polynomial-time algorithm to solve the problem for graphs with pathwidth one. Thus, this paper gives precise analyses of the problem with respect to pathwidth.

Indeed, our reduction for the PSPACE-completeness proof constructs a bipartite series-parallel graph (whose treewidth is two), which is a bipartite planar graph. We note that the problem of finding one k-list coloring of a given graph can be solved in polynomial time for bounded treewidth graphs (and hence for bounded pathwidth graphs) [17]. However, our proof shows that the reconfiguration variant is PSPACE-complete even if treewidth and pathwidth are two. Furthermore, as a byproduct, our reduction can be extended for threshold graphs.

2 Preliminaries

We assume without loss of generality that graphs are simple and connected. Let $G = (V, E)$ be a graph with vertex set V and edge set E; we sometimes denote by $V(G)$ and $E(G)$ the vertex set and edge set of G, respectively. For a vertex v in G, we denote by $d(v)$ the degree of v in G. For a vertex subset $V' \subseteq V$, we denote by $G[V']$ the subgraph of G induced by V'.

We now define the notion of pathwidth [24]. A *path-decomposition* of a graph G is a sequence of subsets X_i of vertices in G such that

(1) $\bigcup_i X_i = V(G)$;
(2) for each $vw \in E(G)$, there is at least one subset X_i with $v, w \in X_i$; and
(3) for any three indices p, q, r such that $p \leq q \leq r$, $X_p \cap X_r \subseteq X_q$.

The *width* of a path-decomposition is defined as $\max_i |X_i| - 1$, and the *pathwidth* of G is the minimum t such that G has a path-decomposition of width t.

To develop our algorithm in Section 4, it is important to notice that every connected graph of pathwidth one is a caterpillar [23]. A caterpillar will be defined in Section 4, but an example can be found in Fig. 4.

For a graph G with a list L, we define the *reconfiguration graph* R_G^L as follows: each node of R_G^L corresponds to a k-list coloring of G, and two nodes of R_G^L are joined by an edge if their corresponding k-list colorings f and f' satisfy $|\{v \in V : f(v) \neq f'(v)\}| = 1$, that is, f' can be obtained from f by changing the color assignment of a single vertex v. We will refer to a *node* of R_G^L in order to distinguish it from a vertex of G. Since we have defined k-list colorings as the nodes of R_G^L, we use graph terms such as adjacency and path for k-list colorings. For notational convenience, we sometimes identify a node of R_G^L with its corresponding k-list coloring of G if it is clear from the context.

Given a graph G with a list L and two k-list colorings f_0 and f_r of G, the k-LIST COLORING RECONFIGURATION problem asks whether the reconfiguration graph R_G^L has a path between the two nodes f_0 and f_r.

3 PSPACE-completeness

In this section, we first prove that the problem is PSPACE-complete even for graphs with pathwidth two. Then, we show in Section 3.3 that the reduction can be extended to proving the PSPACE-completeness of threshold graphs.

A graph is *series-parallel* if it does not contain a subdivision of a complete graph K_4 on four vertices [8]. Note that series-parallel graphs may have super-constant pathwidth, although their treewidth can be bounded by two. We give the following theorem.

Theorem 1. *The* k-LIST COLORING RECONFIGURATION *problem is* PSPACE*-complete even for bipartite series-parallel graphs of pathwidth two.*

It is known that k-LIST COLORING RECONFIGURATION is in PSPACE [5]. Therefore, as a proof of Theorem 1, we give a polynomial-time reduction from the SHORTEST PATH REROUTING problem [4] (defined in Section 3.1) to our problem for bipartite series-parallel graphs of pathwidth two.

3.1 Reconfiguration problem for shortest path

Let H be an unweighted graph, and let s and t be two vertices in H. We call a shortest path in H between s and t simply an *S-path* in H. Note that, since H is unweighted, an S-path indicates a path between s and t having the minimum number of edges. We say that two S-paths P and P' in H are *adjacent* if they differ in exactly one vertex, that is, $|V(P') \setminus V(P)| = 1$ and $|V(P) \setminus V(P')| = 1$ hold. Given two S-paths P_0 and P_r in H, the SHORTEST PATH REROUTING problem asks whether there exists a sequence $\mathcal{P} = \langle P_0, P_1, \ldots, P_\ell \rangle$ of S-paths

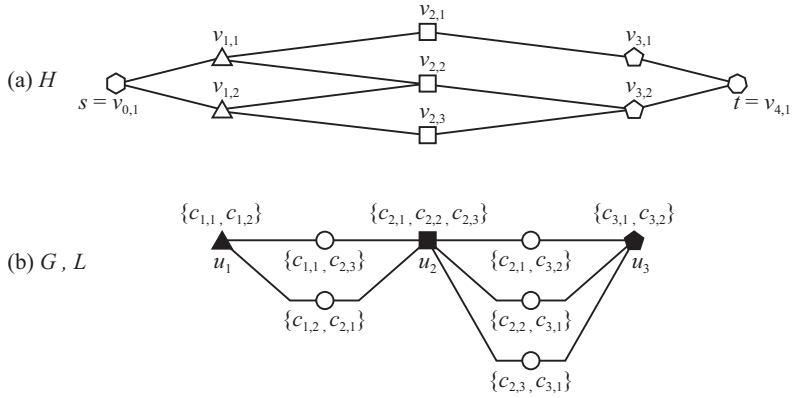

Fig. 2. (a) Graph H for SHORTEST PATH REROUTING, and (b) the corresponding graph G and its list L for k-LIST COLORING RECONFIGURATION.

such that $P_\ell = P_r$, and P_{i-1} and P_i are adjacent for each $i = 1, 2, \ldots, \ell$. This problem is known to be PSPACE-complete [4].

To construct our reduction, we introduce some terms. Let d be the number of edges of an S-path in a graph H. For two vertices v and w in H, we denote by $\mathsf{dist}(v, w)$ the number of edges of a shortest path in H between v and w; then $\mathsf{dist}(s, t) = d$. For each $i \in \{0, 1, \ldots, d\}$, the *layer* $D_{s,i}$ *from* s is defined to be the set of all vertices in H that are placed at distance i from s, that is, $D_{s,i} = \{v \in V(H) : \mathsf{dist}(s, v) = i\}$. Similarly, let $D_{t,j} = \{v \in V(H) : \mathsf{dist}(t, v) = j\}$ for each $j \in \{0, 1, \ldots, d\}$. Then, for each $i \in \{0, 1, \ldots, d\}$, the *layer* D_i *of* H is defined as follows: $D_i = D_{s,i} \cap D_{t,d-i}$. Notice that any S-path in H contains exactly one vertex from each layer D_i, $0 \leq i \leq d$. Then, observe that two S-paths P and P' in H are adjacent if and only if there exists exactly one index $j \in \{1, 2, \ldots, d\}$ such that $V(P) \cap D_j \neq V(P') \cap D_j$. Therefore, we may assume without loss of generality that H consists of only vertices in $\bigcup_{0 \leq i \leq d} D_i$, as illustrated in Fig. 2(a). In the example of Fig. 2(a), all vertices in the same layer D_i are depicted by the same shape, that is, $D_0 = \{s\}$, $D_1 = \{v_{1,1}, v_{1,2}\}$, $D_2 = \{v_{2,1}, v_{2,2}, v_{2,3}\}$, $D_3 = \{v_{3,1}, v_{3,2}\}$ and $D_4 = \{t\}$. Note that both $D_0 = \{s\}$ and $D_d = \{t\}$ always hold.

3.2 Reduction

Given an instance (H, P_0, P_r) of SHORTEST PATH REROUTING, we construct the corresponding instance (G, L, f_0, f_r) of k-LIST COLORING RECONFIGURATION.

Construction of G and L.

We first construct the corresponding graph G with a list L. For each $i \in \{1, 2, \ldots, d-1\}$, let $D_i = \{v_{i,1}, v_{i,2}, \ldots, v_{i,q}\}$; and we introduce a vertex u_i, called a *layer vertex*, to G. The list of each layer vertex u_i is defined as $L(u_i) = \{c_{i,1}, c_{i,2}, \ldots, c_{i,q}\}$, where each color $c_{i,j}$ in $L(u_i)$ corresponds to the vertex $v_{i,j}$ in D_i; assigning color $c_{i,j}$ represents selecting the vertex $v_{i,j}$ as an S-path in H.

We denote by U_{layer} the set of all layer vertices $u_1, u_2, \ldots, u_{d-1}$ in G. In Fig. 2(b), each layer vertex u_i is illustrated as a black vertex which is the same shape as the vertices in D_i.

We then connect layer vertices in G by *forbidden paths* of length two, as follows. Let $v_{i,x} \in D_i$ and $v_{i+1,y} \in D_{i+1}$ be an arbitrary pair of vertices in H such that $v_{i,x}v_{i+1,y} \notin E(H)$. We introduce a vertex w to G, and join w and each of the two layer vertices u_i and u_{i+1} by an edge. The list $L(w)$ of w consists of two colors $c_{i,x}$ and $c_{i+1,y}$ which correspond to the vertices $v_{i,x}$ and $v_{i+1,y}$ in H, respectively. (See the vertices depicted by white circles in Fig. 2(b).) We call such a vertex w in G a $(v_{i,x}, v_{i+1,y})$-*forbidden vertex* or simply a *forbidden vertex*. Notice that there is no proper k-list coloring f such that $f(u_i) = c_{i,x}$ and $f(u_{i+1}) = c_{i+1,y}$; otherwise there is no color in $L(w)$ that can be assigned to w. This property ensures that every k-list coloring of G corresponds to an S-path in H. Let U_{forbid} be the set of all forbidden vertices, then $|U_{\mathsf{forbid}}| = O(|E(\bar{H})|)$ where $\bar{H}$ is the complement graph of H.

This completes the construction of G and L. Since $|V(G)| = O(d+|E(\bar{H})|) = O(|V(H)|^2)$ and $|E(G)| = O(|E(\bar{H})|) = O(|V(H)|^2)$, we can construct G in polynomial time. Clearly, G is a bipartite series-parallel graph of pathwidth two, whose bipartition consists of U_{layer} and U_{forbid}.

Construction of f_0 and f_r.

We now construct two k-list colorings f_0 and f_r of G which correspond to S-paths P_0 and P_r in H, respectively. For each $i \in \{1, 2, \ldots, d-1\}$, let $v_{i,0}$ be the vertex in the layer D_i passed through by the S-path P_0; then we let $f_0(u_i) = c_{i,0}$ for each layer vertex $u_i \in U_{\mathsf{layer}}$. For each $(v_{i,x}, v_{i+1,y})$-forbidden vertex $w \in U_{\mathsf{forbid}}$, we choose an arbitrary color from $c_{i,x}$ and $c_{i+1,y}$ which is assigned to neither u_i nor u_{i+1}. We note that such an available color always exists, because P_0 has an edge between the two vertices corresponding to the colors $f_0(u_i)$ and $f_0(u_{i+1})$, and hence at least one of $f_0(u_i) \neq c_{i,x}$ and $f_0(u_{i+1}) \neq c_{i+1,y}$ holds for each $(v_{i,x}, v_{i+1,y})$-forbidden vertex. Similarly, we construct f_r. This completes the construction of the corresponding instance (G, L, f_0, f_r).

Correctness of the reduction.

To show the correctness of this reduction, we give the following lemma.

Lemma 1. *(H, P_0, P_r) is a yes-instance if and only if (G, L, f_0, f_r) is a yes-instance.*

Proof. We first prove the if-part. Suppose that the reconfiguration graph R_G^L has a path between the two nodes f_0 and f_r. We can classify the recoloring steps into the following two types: (1) recoloring a layer vertex in U_{layer}, and (2) recoloring a forbidden vertex in U_{forbid}. Therefore, the path in R_G^L can be divided into sub-paths, intermittently at each edge corresponding to a recoloring step of type (1) above; all edges in a sub-path correspond to recoloring steps of type (2) above. Therefore, all nodes in each sub-path correspond to the same S-path in H. Furthermore, any two consecutive sub-paths correspond to two adjacent S-paths (that differ in only one vertex), because the sub-paths are divided by a

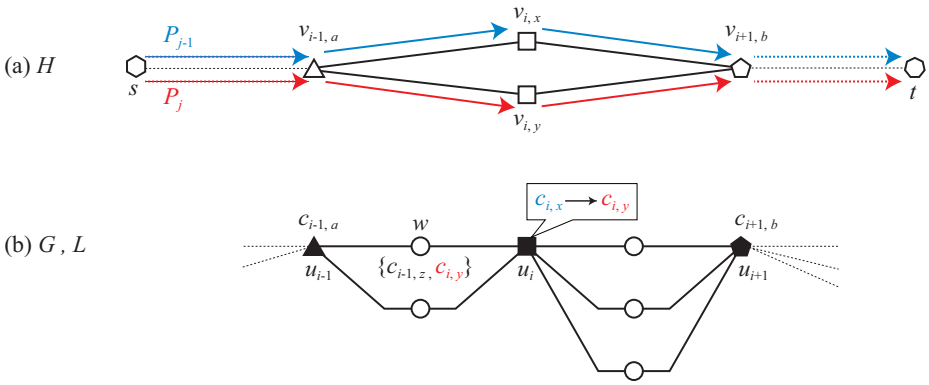

Fig. 3. (a) Two adjacent S-paths P_{j-1} (blue) and P_j (red), and (b) the corresponding recoloring steps.

recoloring step of type (1). Thus, we can construct a sequence of adjacent S-paths that transforms P_0 into P_r.

We then prove the only-if-part. Suppose that there exists a sequence $\mathcal{P} = \langle P_0, P_1, \ldots, P_\ell \rangle$ of S-paths such that $P_\ell = P_r$, and P_{j-1} and P_j are adjacent for each $j = 1, 2, \ldots, \ell$. For two adjacent S-paths P_{j-1} and P_j, $j \in \{1, 2, \ldots, \ell\}$, assume that P_j is obtained from P_{j-1} by replacing $v_{i,x} \in V(P_{j-1}) \cap D_i$ with $v_{i,y} \in V(P_j) \cap D_i$. (See Fig. 3(a).) This rerouting step from $v_{i,x}$ to $v_{i,y}$ corresponds to recoloring the layer vertex $u_i \in U_{\mathsf{layer}}$ from the color $c_{i,x}$ to $c_{i,y}$. (See Fig. 3(b).) To do so, if there is a forbidden vertex $w \in U_{\mathsf{forbid}}$ which is adjacent with u_i and receives the color $c_{i,y}$, we first need to recolor w from $c_{i,y}$ to another color $c_{q,z} \in L(w) \setminus \{c_{i,y}\}$, where $q \in \{i-1, i+1\}$. (Note that, since P_{j-1} corresponds to a feasible k-list coloring of G, no forbidden vertex adjacent with u_i receives the color $c_{i,x}$.) Since w is a forbidden vertex, we know that $|L(w)| = 2$. Let $v_{q,a}$ be the vertex in the layer D_q which is adjacent with $v_{i,x}$ and $v_{i,y}$ in P_{j-1} and P_j, respectively; and hence the layer vertex u_q receives the color $c_{q,a}$. (Figure 3 illustrates the case where $q = i - 1$.) Then, since each forbidden vertex is placed only for a pair of vertices in H which is *not* joined by an edge, we observe that the color $c_{q,z} \in L(w) \setminus \{c_{i,y}\}$ is different from $c_{q,a}$. Therefore, we can recolor w from $c_{i,y}$ to another color $c_{q,z}$ without recoloring any other vertices. Since U_{forbid} forms an independent set of G, we can apply this recoloring step to all forbidden vertices independently. Now, any neighbor of u_i is colored with neither $c_{i,x}$ nor $c_{i,y}$, and hence we can recolor u_i from $c_{i,x}$ to $c_{i,y}$. Thus, R_G^L has a path between f_0 and f_r which corresponds to $\mathcal{P}$. □

3.3 Threshold graphs

In this subsection, we extend our reduction in Section 3.2 to threshold graphs. A graph G is *threshold* if there exist a real number s and a mapping $\omega : V(G) \to \mathbb{R}$ such that $xy \in E(G)$ if and only if $\omega(x) + \omega(y) \geq s$ [8], where $\mathbb{R}$ is the set of all real numbers.

Theorem 2. *The k-*LIST COLORING RECONFIGURATION *problem is* PSPACE*-complete even for threshold graphs.*

Proof. We modify the graph G constructed in Section 3.2, as follows: join all pairs of vertices in U_{layer}, and join each vertex in U_{forbid} with all vertices in U_{layer}. Let G' be the resulting graph, then $G[U_{\mathsf{layer}}]$ forms a clique and $G[U_{\mathsf{forbid}}]$ forms an independent set. Notice that G' is a threshold graph; set the threshold $s = 1$, and the mapping $\omega(v) = 1$ for each layer vertex $v \in U_{\mathsf{layer}}$ and $\omega(u) = 0$ for each forbidden vertex $u \in U_{\mathsf{forbid}}$.

Consider any pair of vertices u and v such that $uv \in E(G') \setminus E(G)$, that is, they are joined by a new edge for constructing G' from G. Then, by the construction in Section 3.2, the two lists $L(u)$ and $L(v)$ contain no color in common. Therefore, adding new edges to G does not affect the existence of k-list coloring: more formally, any k-list coloring of G is a k-list coloring of G', and vice versa. Thus, by Lemma 1, an instance (H, P_0, P_r) of SHORTEST PATH REROUTING is a yes-instance if and only if (G', L, f_0, f_r) is a yes-instance. □

4 Algorithm for Graphs with Pathwidth One

In contrast to Theorem 1, we give the following theorem in this section.

Theorem 3. *The k-*LIST COLORING RECONFIGURATION *problem can be solved in polynomial time for graphs with pathwidth one.*

As a proof of Theorem 3, we give such an algorithm. However, since every connected graph of pathwidth one is a caterpillar [23], it suffices to develop a polynomial-time algorithm for caterpillars.

A *caterpillar* G is a tree whose vertex set $V(G)$ can be partitioned into two subset V_{S} and V_{L} such that $G[V_{\mathrm{S}}]$ forms a path and each vertex in V_{L} is incident to exactly one vertex in V_{S}. We may assume without loss of generality that the two endpoints of the path $G[V_{\mathrm{S}}]$ are of degree one in the whole graph G. (See v_1 and v_{10} in Fig. 4.) We call each vertex in V_{S} a *spine vertex* of G, and each vertex in V_{L} a *leaf* of G.

We assume that all vertices in G are ordered as $v_1, v_2, \ldots, v_n$ by the breadth-first search starting with the endpoint (degree-1 vertex) of the spine path $G[V_{\mathrm{S}}]$ with the priority to leaves; that is, when we visit a spine vertex v, we first visit all leaves of v and then visit the unvisited spine vertex. (See Fig. 4 for example.)

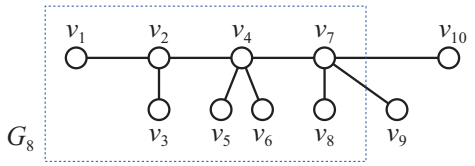

Fig. 4. A caterpillar G and its vertex ordering, where the subgraph surrounded by a dotted rectangle corresponds to G_8.

For each index $i \in \{1, 2, \ldots, n\}$, we let $V_i = \{v_1, v_2, \ldots, v_i\}$ and $G_i = G[V_i]$. Then, clearly $G_n = G$. For each index $i \in \{1, 2, \ldots, n\}$, let $\mathsf{sp}(i)$ be the latest spine vertex in V_i, that is, $\mathsf{sp}(i) = v_i$ if v_i is a spine vertex, otherwise $\mathsf{sp}(i)$ is the unique neighbor of v_i. Then, v_i is adjacent with only the spine vertex $\mathsf{sp}(i-1)$ in G_i for each $i \in \{2, 3, \ldots, n\}$.

We can restrict the length of each list without loss of generality, as follows. Note that the following lemma holds for any graph.

Lemma 2. *For an instance (G', L', f'_0, f'_r), one can obtain another instance (G, L, f_0, f_r) in polynomial time such that $2 \leq |L(v)| \leq d(v) + 1$ for each vertex $v \in V(G)$, and (G', L', f'_0, f'_r) is a yes-instance if and only if (G, L, f_0, f_r) is a yes-instance.*

Proof. If $|L(v)| = 1$ for a vertex $v \in V(G')$, then any k-list coloring f' of G' assigns the same color $c \in L(v)$ to v. Therefore, f' never assigns c to any neighbor u of v. We can thus delete v from G' and set $L(u) := L(u) \setminus \{c\}$ for all neighbors u of v in G'. Clearly, this modification does not affect the reconfigurability (i.e., the existence or non-existence of a path in the reconfiguration graph).

If $|L(v)| \geq d(v) + 2$ for a vertex $v \in V(G')$, we simply delete v from G' without any modification of lists; let G be the resulting graph. Let f be any k-list coloring of G, and consider a recoloring step for a neighbor u of v from the current color $c = f(u)$ to another color c'. If c' is not assigned to v in G', we can directly recolor u from c to c'. Thus, suppose that c' is assigned to v in G'. Then, since $|L(v)| \geq d(v) + 2$, there is at least one color $c^* \in L(v)$ which is not c' and is not assigned to any of $d(v)$ neighbors of v by f. Therefore, we first recolor v from c' to c^*, and then recolor u from c to c'. In this way, any recoloring step in G can be simulated in G', and hence the modification does not affect the reconfigurability.

Thus, we can obtain an instance such that $2 \leq |L(v)| \leq d(v)+1$ holds for each vertex v without affecting the reconfigurability. Clearly, the modified instance can be constructed in polynomial time. □

Therefore, in the remainder of this section, we assume that G is a (connected) caterpillar and $2 \leq |L(v)| \leq d(v) + 1$ holds for every vertex $v \in V(G)$. In particular, $|L(v)| = 2$ for every leaf v of G.

4.1 Idea and definitions

The main idea of our algorithm is to extend techniques developed for SHORTEST PATH REROUTING [3], and apply them to k-LIST COLORING RECONFIGURATION for caterpillars. Our algorithm employs a dynamic programming method based on the vertex ordering $v_1, v_2, \ldots, v_n$ of G.

For each $i \in \{1, 2, \ldots, n\}$, let $\mathsf{R}^L_{G_i}$ be the reconfiguration graph for the subgraph G_i and the list L. Then, $\mathsf{R}^L_{G_i}$ contains all k-list colorings of G_i as its nodes. Our algorithm efficiently constructs $\mathsf{R}^L_{G_i}$ for each $i = 1, 2, \ldots, n$, in this order. However, of course, the number of nodes in $\mathsf{R}^L_{G_i}$ cannot be bounded by a

polynomial size in general. We thus use the property that the vertex v_{i+1} (will be added to G_i) is adjacent with only the spine vertex $\mathsf{sp}(i)$ in G_{i+1}; and we "encode" the reconfiguration graph $\mathsf{R}^L_{G_i}$ into a polynomial size with keeping the information of (1) the color assigned to $\mathsf{sp}(i)$, and (2) the connectivity of nodes in $\mathsf{R}^L_{G_i}$.

Before explaining the encoding methods, we first note that it suffices to focus on only one connected component in $\mathsf{R}^L_{G_i}$ which contains the restriction of f_0, where the *restriction* of a k-list coloring f of a graph G to a subgraph G' is a k-list coloring g of G' such that $g(v) = f(v)$ hold for all vertices $v \in V(G')$. For notational convenience, we denote by $f[V_i]$ the restriction of a k-list coloring f of a caterpillar G to its subgraph G_i. Then, we have the following lemma.

Lemma 3. *Let g be a k-list coloring of G_i such that $f_0[V_i]$ and g are contained in the same connected component in $\mathsf{R}^L_{G_i}$. Then, for each $j \in \{1, 2, \ldots, i-1\}$, $f_0[V_j]$ and $g[V_j]$ are contained in the same connected component in $\mathsf{R}^L_{G_j}$.*

Proof. Assume that $f_0[V_i]$ and g are contained in the same connected component in $\mathsf{R}^L_{G_i}$. Then, there exists a path in $\mathsf{R}^L_{G_i}$ between $f_0[V_i]$ and g. We contract all edges in the path that correspond to recoloring vertices in $V_i \setminus V_j$. Since each edge in the resulting path corresponds to recoloring only one vertex in V_j, the resulting path must be contained as a path in $\mathsf{R}^L_{G_j}$ between $f_0[V_j]$ and $g[V_j]$. Therefore $f_0[V_j]$ and $g[V_j]$ are contained in the same connected component in $\mathsf{R}^L_{G_j}$, and hence the lemma follows. □

From now on, we thus focus on only the connected component of $\mathsf{R}^L_{G_i}$ which contains $f_0[V_i]$. Since the list is fixed to be L in the remainder of this section, we simply denote by R_i the reconfiguration graph $\mathsf{R}^L_{G_i}$, and by R^0_i the connected component of $\mathsf{R}_i = \mathsf{R}^L_{G_i}$ containing $f_0[V_i]$.

Encoding graph.

We now partition the nodes of R^0_i into several subsets with respect to (1) the color assigned to $\mathsf{sp}(i)$, and (2) the connectivity of nodes in R^0_i. For two nodes g and g' of R^0_i with $g(\mathsf{sp}(i)) = g'(\mathsf{sp}(i))$, we write $g \sim_{\mathsf{sp}(i)} g'$ if g can be reconfigured into g' without recoloring the color assigned to the vertex $\mathsf{sp}(i)$, that is, R^0_i has a path $\langle g_1, g_2, \ldots, g_\ell \rangle$ such that $g_1 = g$, $g_\ell = g'$, and $g_j(\mathsf{sp}(i)) = g(\mathsf{sp}(i)) = g'(\mathsf{sp}(i))$ holds for every $j \in \{1, 2, \ldots, \ell\}$. Since the adjacency relation on k-list colorings is symmetric (i.e., R_i is an undirected graph), it is easy to see that $\sim_{\mathsf{sp}(i)}$ is an equivalence relation. Thus, the node set of R^0_i can be uniquely partitioned by the relation $\sim_{\mathsf{sp}(i)}$. We denote by $\mathcal{G}^0_i$ the partition of the node set of R^0_i into equivalence classes with respect to $\sim_{\mathsf{sp}(i)}$.

We finally define our dynamic programming table. For each subgraph G_i, $i \in \{1, 2, \ldots, n\}$, our algorithm keeps track of four information $(H_i, \mathsf{col}_i, \mathsf{ini}_i, \mathsf{tar}_i)$, defined as follows.

- The *encoding graph* H_i of R^0_i which can be obtained from R^0_i by contracting each node set in $\mathcal{G}^0_i$ into a single node. (See Fig. 5 as an example.) We will refer to an *e-node* of H_i in order to distinguish it from a node of R^0_i. (Thus,

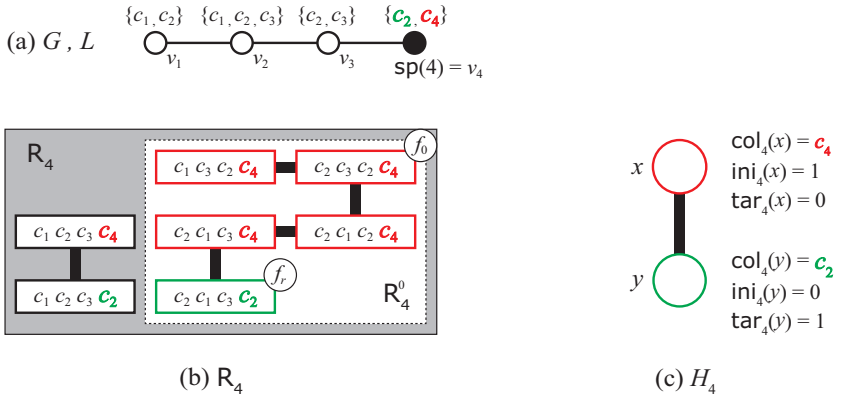

Fig. 5. (a) A caterpillar $G = G_4$, (b) the reconfiguration graph R_4 consisting of all k-list colorings of G, and (c) the encoding graph H_4 of R_4^0 consisting of two e-nodes x and y, where each k-list coloring in (b) is represented as the sequence of colors assigned to the vertices in G from left to right.

each node refers to a k-list coloring of G_i, and each e-node refers to a set of k-list colorings of G_i.) For each e-node $x \in V(H_i)$, we denote by $\Phi_i(x)$ the set of all nodes in R_i^0 that were contracted into x. (Note that we do not compute $\Phi_i(x)$, but use it only for definitions and proofs.)

- The color $\mathsf{col}_i(x) \in L(\mathsf{sp}(i))$ for each e-node $x \in V(H_i)$, which is assigned to $\mathsf{sp}(i)$ in common by the nodes (i.e., k-list colorings of G_i) in $\Phi_i(x)$.
- The label $\mathsf{ini}_i(x) \in \{0, 1\}$ for each e-node $x \in V(H_i)$, such that $\mathsf{ini}_i(x) = 1$ if $f_0[V_i] \in \Phi_i(x)$, otherwise $\mathsf{ini}_i(x) = 0$.
- The label $\mathsf{tar}_i(x) \in \{0, 1\}$ for each e-node $x \in V(H_i)$, such that $\mathsf{tar}_i(x) = 1$ if $f_r[V_i] \in \Phi_i(x)$, otherwise $\mathsf{tar}_i(x) = 0$.

To prove Theorem 3, we give a polynomial-time algorithm which computes $(H_i, \mathsf{col}_i, \mathsf{ini}_i, \mathsf{tar}_i)$ for each subgraph G_i, $i \in \{1, 2, \ldots, n\}$, by means of dynamic programming. Then, the problem can be solved as in the following lemma.

Lemma 4. *(G, L, f_0, f_r) is a yes-instance if and only if the encoding graph H_n contains a node x such that $\mathsf{tar}_n(x) = 1$.*

Proof. Since H_n contains a node x such that $\mathsf{tar}_n(x) = 1$, we have $f_r[V_n] = f_r \in \Phi_n(x)$. Recall that H_n is the encoding graph of R_n^0 which contains the k-list coloring f_0 of G as a node. Since $\Phi_n(x) \subseteq V(\mathsf{R}_n^0)$, the lemma clearly follows. □

4.2 Algorithm

As the initialization, we first consider the case where $i = 1$, that is, we compute $(H_1, \mathsf{col}_1, \mathsf{ini}_1, \mathsf{tar}_1)$. (See Fig. 6(d) as an example.) Note that G_1 consists of a

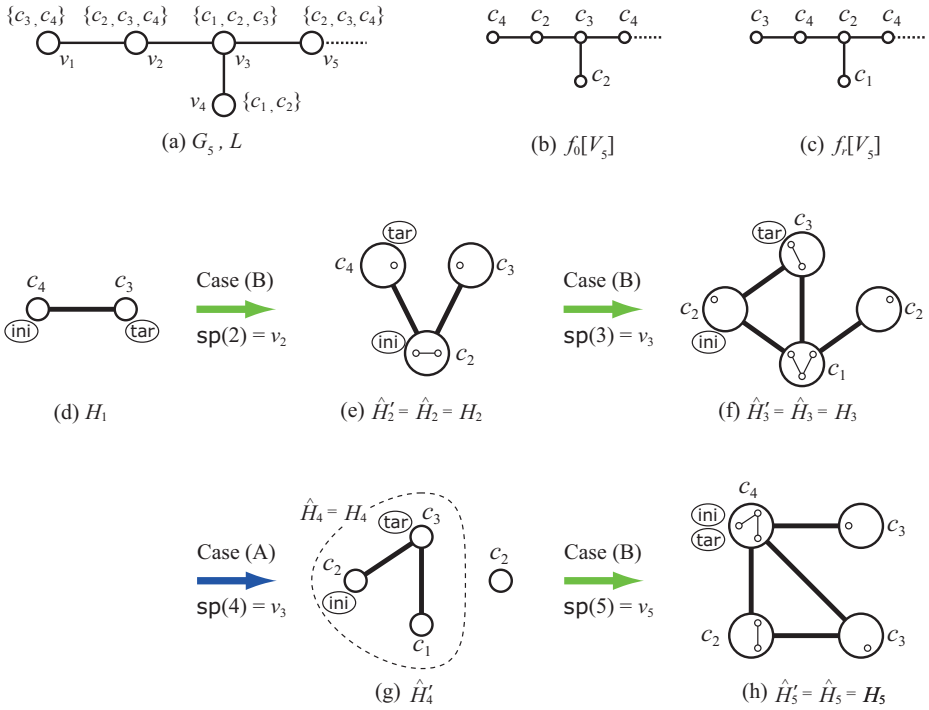

Fig. 6. Application of our algorithm. In (d)–(h), $\mathsf{col}_i(x) \in L(\mathsf{sp}(i))$ is attached to each e-node x, and the e-nodes x with $\mathsf{ini}_i(x) = 1$ and $\mathsf{tar}_i(x) = 1$ have the labels "ini" and "tar," respectively. Furthermore, in (e), (f) and (h), the small graph contained in each e-node x of H_i represents the subgraph of H_{i-1} induced by $\mathsf{EN}(x)$.

single vertex v_1, and recall that v_1 is a spine vertex of degree one. By Lemma 2 we then have $|L(v_1)| = 2$. Therefore, the reconfiguration graph R_1 is a complete graph on $|L(v_1)| = 2$ nodes such that each node corresponds to a k-list coloring of G_1 which assigns a distinct color to the vertex $\mathsf{sp}(1) = v_1$. Since R_1 is complete and contains the node $f_0[V_1]$, we have $\mathsf{R}_1^0 = \mathsf{R}_1$. Furthermore, $H_1 = \mathsf{R}_1^0$ since all nodes in R_1^0 assign distinct colors in $L(v_1)$ to $\mathsf{sp}(1) = v_1$. Then, for each e-node x of H_1 corresponding to the set consisting of a single k-list coloring g of G_1, we set

$$\mathsf{col}_1(x) = g(v_1);$$
$$\mathsf{ini}_1(x) = \begin{cases} 1 & \text{if } g(v_1) = f_0(v_1), \\ 0 & \text{otherwise;} \end{cases}$$
$$\mathsf{tar}_1(x) = \begin{cases} 1 & \text{if } g(v_1) = f_r(v_1), \\ 0 & \text{otherwise.} \end{cases}$$

For $i \geq 2$, suppose that we have already computed $(H_{i-1}, \mathsf{col}_{i-1}, \mathsf{ini}_{i-1}, \mathsf{tar}_{i-1})$. Then, we compute $(H_i, \mathsf{col}_i, \mathsf{ini}_i, \mathsf{tar}_i)$, as follows.

Fig. 7. The graph G_i for (a) $v_i \in V_{\mathrm{L}}$ and (b) $v_i \in V_{\mathrm{S}}$.

Case (A): v_i is a leaf in V_{L}. (See Figs. 6(g) and 7(a).)

By Lemma 2 we have $|L(v_i)| = 2$ in this case; let $L(v_i) = \{c_1, c_2\}$. Recall that v_i is adjacent with only the spine vertex $\mathsf{sp}(i-1)$ in G_i. Furthermore, $\mathsf{sp}(i) = \mathsf{sp}(i-1)$ in this case.

Let $H_{i-1}^{c_1}$ be the subgraph of H_{i-1} obtained by deleting all e-nodes y in H_{i-1} with $\mathsf{col}_{i-1}(y) = c_1$. Then, $H_{i-1}^{c_1}$ encodes all nodes of R_{i-1}^0 that do not assign the color c_1 to $\mathsf{sp}(i-1)$. Thus, we can extend each k-list coloring h of G_{i-1} encoded in $H_{i-1}^{c_1}$ to a k-list coloring g of G_i such that $g(v_i) = c_1$ and $g(v) = h(v)$ for all vertices $v \in V_{i-1}$. Similarly, let $H_{i-1}^{c_2}$ be the subgraph of H_{i-1} obtained by deleting all e-nodes z in H_{i-1} with $\mathsf{col}_{i-1}(z) = c_2$.

We define an encoding graph $\hat{H}_i'$ as $V(\hat{H}_i') = V(H_{i-1}^{c_1}) \cup V(H_{i-1}^{c_2})$ and $E(\hat{H}_i') = E(H_{i-1}^{c_1}) \cup E(H_{i-1}^{c_2})$; and let $\hat{H}_i$ be the connected component of $\hat{H}_i'$ that contains the e-node x such that $\mathsf{ini}_{i-1}(x) = 1$. For each e-node x in $\hat{H}_i$, let $\hat{\mathsf{col}}_i(x) = \mathsf{col}_{i-1}(x)$, $\hat{\mathsf{ini}}_i(x) = \mathsf{ini}_{i-1}(x)$ and $\hat{\mathsf{tar}}_i(x) = \mathsf{tar}_{i-1}(x)$. Then, we have the following lemma.

Lemma 5. *For a leaf $v_i \in V_{\mathrm{L}}$, $(H_i, \mathsf{col}_i, \mathsf{ini}_i, \mathsf{tar}_i) = (\hat{H}_i, \hat{\mathsf{col}}_i, \hat{\mathsf{ini}}_i, \hat{\mathsf{tar}}_i)$.*

Case (B): v_i is a spine vertex in V_{S}. (See Figs. 6(e), (f), (h) and 7(b).)

In this case, notice that $\mathsf{sp}(i) = v_i$ in G_i, and hence we need to update col_i according to the color assigned to v_i.

We first define an encoding graph $\hat{H}_i'$, as follows. For a color $c \in L(v_i)$, let H_{i-1}^c be the subgraph of H_{i-1} obtained by deleting all e-nodes y in H_{i-1} with $\mathsf{col}_{i-1}(y) = c$. For each connected component in H_{i-1}^c, we add a new e-node x to $\hat{H}_i'$ such that $\hat{\mathsf{col}}_i(x) = c$; we denote by $\mathsf{EN}(x)$ the set of all e-nodes in H_{i-1}^c that correspond to x. We apply this operation to all colors in $L(v_i)$. We then add edges to $\hat{H}_i'$: two e-nodes x and y in $\hat{H}_i'$ are joined by an edge if and only if $\mathsf{EN}(x) \cap \mathsf{EN}(y) \neq \emptyset$.

We now define $\hat{\mathsf{ini}}_i(x)$ and $\hat{\mathsf{tar}}_i(x)$ for each e-node x in $\hat{H}_i'$, as follows:

$$\hat{\mathsf{ini}}_i(x) = \begin{cases} 1 & \text{if } \hat{\mathsf{col}}_i(x) = f_0(v_i) \text{ and} \\ & \mathsf{EN}(x) \text{ contains an e-node } y \text{ with } \mathsf{ini}_{i-1}(y) = 1; \\ 0 & \text{otherwise,} \end{cases}$$

and

$$\hat{\mathsf{tar}}_i(x) = \begin{cases} 1 & \text{if } \hat{\mathsf{col}}_i(x) = f_r(v_i) \text{ and} \\ & \mathsf{EN}(x) \text{ contains an e-node } y \text{ with } \mathsf{tar}_{i-1}(y) = 1; \\ 0 & \text{otherwise.} \end{cases}$$

Let $\hat{H}_i$ be the connected component of $\hat{H}'_i$ that contains the e-node x such that $\hat{\mathsf{ini}}_i(x) = 1$. Then, we have the following lemma.

Lemma 6. *For a spine vertex $v_i \in V_{\mathrm{S}}$, $(H_i, \mathsf{col}_i, \mathsf{ini}_i, \mathsf{tar}_i) = (\hat{H}_i, \hat{\mathsf{col}}_i, \hat{\mathsf{ini}}_i, \hat{\mathsf{tar}}_i)$.*

4.3 Running time

We now estimate the running time of our algorithm in Section 4.2. The following is the key lemma for the estimation.

Lemma 7. *For each index $i \in \{1, 2, \ldots, n\}$,*

$$|V(H_i)| \leq \begin{cases} 2 & \text{if } i = 1; \\ |V(H_{i-1})| + d(v_i) & \text{otherwise.} \end{cases}$$

In particular, $|V(H_n)| = O(n)$, where n is the number of vertices in G.

By Lemma 7 each encoding graph H_i is of size $O(n)$ for each $i \in \{1, 2, \ldots, n\}$. Therefore, our algorithm clearly runs in polynomial time.

This completes the proof of Theorem 3. □

5 Concluding Remarks

In this paper, we gave precise analyses of the k-list coloring reconfiguration problem with respect to pathwidth: the problem is solvable in polynomial time for graphs with pathwidth one, while it is PSPACE-complete for graphs with pathwidth two.

Very recently, Wrochna [25] gave another proof for the PSPACE-completeness of k-list coloring reconfiguration for graphs with pathwidth two. His reduction is constructed from a PSPACE-complete problem, called H-Word Reachability.

Acknowledgments

We are grateful to Daichi Fukase and Yuma Tamura for fruitful discussions with them. This work is partially supported by JSPS KAKENHI Grant Numbers 25106504 and 25330003.

References

1. Bonamy, M., Johnson, M., Lignos, I., Patel, V., Paulusma, D.: Reconfiguration graphs for vertex colourings of chordal and chordal bipartite graphs. J. Combinatorial Optimization 27, pp. 132–143 (2014)
2. Bonamy, M., Bousquet, N.: Recoloring bounded treewidth graphs. Electronic Notes in Discrete Mathematics 44, pp. 257–262 (2013)
3. Bonsma, P.: Rerouting shortest paths in planar graphs. Proc. of FSTTCS 2012, pp. 337–349 (2012)

4. Bonsma, P.: The complexity of rerouting shortest paths. Theoretical Computer Science 510, pp. 1–12 (2013)
5. Bonsma, P., Cereceda, L.: Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. Theoretical Computer Science 410, pp. 5215–5226 (2009)
6. Bonsma, P., Kamiński, M., Wrochna, M.: Reconfiguring independent sets in claw-free graphs. Proc. of SWAT 2014, LNCS 8503, pp. 86–97 (2014)
7. Bonsma, P., Mouawad, A.E.: The complexity of bounded length graph recoloring. arXiv:1404.0337 (2014)
8. Brandstädt, A., Le, V.B., Spinrad, J.P.: Graph Classes: A Survey. SIAM (1999)
9. Cereceda, L.: Mixing Graph Colourings. Ph.D. Thesis, London School of Economics and Political Science (2007)
10. Cereceda, L., van den Heuvel, J., Johnson, M.: Finding paths between 3-colorings. J. Graph Theory 67, pp. 69–82 (2011)
11. Gopalan, P., Kolaitis, P.G., Maneva, E.N., Papadimitriou, C.H.: The connectivity of Boolean satisfiability: computational and structural dichotomies. SIAM J. Computing 38, pp. 2330–2355 (2009)
12. Hearn, R.A., Demaine, E.D.: PSPACE-completeness of sliding-block puzzles and other problems through the nondeterministic constraint logic model of computation. Theoretical Computer Science 343, pp. 72–96 (2005)
13. Ito, T., Demaine, E.D., Harvey, N.J.A., Papadimitriou, C.H., Sideri, M., Uehara, R., Uno, Y.: On the complexity of reconfiguration problems. Theoretical Computer Science 412, pp. 1054–1065 (2011)
14. Ito, T., Kamiński, M., Demaine, E.D.: Reconfiguration of list edge-colorings in a graph. Discrete Applied Mathematics 160, pp. 2199–2207 (2012)
15. Ito, T., Kawamura, K., Ono, H., Zhou, X.: Reconfiguration of list $L(2,1)$-labelings in a graph. To appear in Theoretical Computer Science. DOI: 10.1016/j.tcs.2014.04.011
16. Ito, T., Kawamura, K., Zhou, X.: An improved sufficient condition for reconfiguration of list edge-colorings in a tree. IEICE Trans. on Information and Systems E95-D, pp. 737–745 (2012)
17. Jansen, K., Scheffler, P.: Generalized coloring for tree-like graphs. Discrete Applied Mathematics 75, pp. 135–155 (1997)
18. Johnson, M., Kratsch, D., Kratsch, S., Patel, V., Paulusma, D.: Colouring reconfiguration is fixed-parameter tractable. arXiv:1403.6347 (2014)
19. Kamiński, M., Medvedev, P., Milanič, M.: Shortest paths between shortest paths. Theoretical Computer Science 412, pp. 5205–5210 (2011)
20. Kamiński, M., Medvedev, P., Milanič, M.: Complexity of independent set reconfigurability problems. Theoretical Computer Science 439, pp. 9–15 (2012)
21. Makino, K., Tamaki, S., Yamamoto, M.: An exact algorithm for the Boolean connectivity problem for k-CNF. Theoretical Computer Science Theor. Comp. Sci. 412, pp. 4613–4618 (2011)
22. Mouawad, A.E., Nishimura, N., Raman, V., Simjour, N., Suzuki, A.: On the parameterized complexity of reconfiguration problems. IPEC 2013, pp. 281–294 (2013)
23. Proskurowski, A., Telle, J.A.: Classes of graphs with restricted interval models. Discrete Mathematics & Theoretical Computer Science 3, pp. 167–176 (1999)
24. Robertson, N., Seymour, P.: Graph minors. I. Excluding a forest. J. Combinatorial Theory, Series B, 35, pp. 39–61 (1983)
25. Wrochna, M.: Reconfiguration in bounded bandwidth and treedepth. arXiv:1405.0847 (2014)

A Proofs Omitted from Section 4

We first introduce some notation.

For $i \geq 2$, let h be any node (i.e., a k-list coloring of G_i) in the reconfiguration graph R_{i-1}. Recall that the vertex v_i is adjacent with only the spine vertex $\mathsf{sp}(i-1)$ in G_i. Let c be any color in $L(v_i) \setminus \{f(\mathsf{sp}(i-1))\}$. Then, we say that h can be *extended by* c to a k-list coloring g of G_i such that $g(v_i) = c$ and $g(v) = h(v)$ for all vertices $v \in V_{i-1}$; we simply denote such an extension by $h + c = g$. Note that g is a node in R_i.

For any e-node y in the encoding graph H_{i-1} of R^0_{i-1}, recall that $\Phi_{i-1}(y)$ is the set of all nodes in R^0_{i-1} that were contracted into y. For a color $c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(y)\}$, let $\Phi_{i-1}(y) \oplus c = \{h + c : h \in \Phi_{i-1}(y)\}$, that is, $\Phi_{i-1}(y) \oplus c$ is the set of nodes in R_i that are extended by c from nodes in $\Phi_{i-1}(y)$ $\big(\subseteq V(\mathsf{R}^0_{i-1})\big)$.

A.1 Proof of Lemma 5

In this subsection, we prove Lemma 5. Therefore, suppose that we have already computed $(H_{i-1}, \mathsf{col}_{i-1}, \mathsf{ini}_{i-1}, \mathsf{tar}_{i-1})$ for $i \geq 2$, and assume that $v_i \in V_{\mathrm{L}}$ and $L(v_i) = \{c_1, c_2\}$.

We first prove that $V(H_i) \subseteq V(\hat{H}'_i)$ holds, as in the following lemma.

Lemma 8. *Let x be any e-node in H_i. Then, there exists an e-node $\hat{x}$ in $V(\hat{H}'_i) = V(H^{c_1}_{i-1}) \cup V(H^{c_2}_{i-1})$ such that $\bigcup\big\{\Phi_{i-1}(\hat{x}) \oplus c : c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}\big\} = \Phi_i(x)$.*

Proof. Let g be any node in $\Phi_i(x) \subseteq V(\mathsf{R}^0_i)$. Then, by Lemma 3 the node $g[V_{i-1}]$ is contained in R^0_{i-1}, and hence there exists an e-node $\hat{x}$ in H_{i-1} such that $g[V_{i-1}] \in \Phi_{i-1}(\hat{x})$. Notice that, since $L(v_i) = \{c_1, c_2\}$, we have $V(H^{c_1}_{i-1}) \cup V(H^{c_2}_{i-1}) = V(H_{i-1})$. Thus, the e-node $\hat{x}$ is contained in $\hat{H}'_i$, too. Therefore, we prove that $\bigcup\big\{\Phi_{i-1}(\hat{x}) \oplus c : c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}\big\} = \Phi_i(x)$.

Let g' be any node in $\Phi_i(x)$. Then, $g \sim_{\mathsf{sp}(i)} g'$, and hence R^0_i contains a path $\langle g_1, g_2, \ldots, g_\ell \rangle$ such that $g_1 = g$, $g_\ell = g'$ and $g_j(\mathsf{sp}(i)) = g(\mathsf{sp}(i)) = g'(\mathsf{sp}(i))$ for all $j \in \{1, 2, \ldots, \ell\}$. Thus, we can obtain a path $\langle g_1[V_{i-1}], g_2[V_{i-1}], \ldots, g_\ell[V_{i-1}] \rangle$ in which the spine vertex $\mathsf{sp}(i)$ always receives the same color $g(\mathsf{sp}(i)) = g'(\mathsf{sp}(i))$; note that $g_j[V_{i-1}] = g_{j+1}[V_{i-1}]$ may hold if v_i is recolored, but we can simply drop $g_{j+1}[V_{i-1}]$ in such a case. Since $\mathsf{sp}(i) = \mathsf{sp}(i-1)$ for the case where $v_i \in V_{\mathrm{L}}$, we have $g[V_{i-1}] \sim_{\mathsf{sp}(i-1)} g'[V_{i-1}]$. Therefore, $g'[V_{i-1}] \in \Phi_{i-1}(\hat{x})$. Since $\mathsf{col}_{i-1}(\hat{x})$ represents the color assigned to $\mathsf{sp}(i) = \mathsf{sp}(i-1)$ and v_i is adjacent with $\mathsf{sp}(i-1)$ in G_i, the color $g'(v_i)$ is clearly contained in $L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}$. We thus have $g'[V_{i-1}] + g'(v_i) = g' \in \bigcup\big\{\Phi_{i-1}(\hat{x}) \oplus c : c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}\big\}$.

Let g'' be any node in $\bigcup\big\{\Phi_{i-1}(\hat{x}) \oplus c : c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}\big\}$ such that $g'' = h + c$ for some node h in $\Phi_{i-1}(\hat{x})$ and $c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}$. Since $h \in \Phi_{i-1}(\hat{x})$ and $g[V_{i-1}] \in \Phi_{i-1}(\hat{x})$, we have $g[V_{i-1}] \sim_{\mathsf{sp}(i-1)} h$ and hence R^0_{i-1} contains a path $\langle h_1, h_2, \ldots, h_\ell \rangle$ such that $h_1 = g[V_{i-1}]$, $g_\ell = h$ and $h_j(\mathsf{sp}(i-1)) = h(\mathsf{sp}(i-1))$ for all $j \in \{1, 2, \ldots, \ell\}$. Since $\mathsf{sp}(i) = \mathsf{sp}(i-1)$ and $g(v_i) \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}$, the sequence $\langle h_1 + g(v_i), h_2 + g(v_i), \ldots, h_\ell + g(v_i) \rangle$ is a path in R^0_i. If $g(v_i) \neq c$, we add one more adjacent node $h_\ell + c$ to the last. Since $h_1 + g(v_i) = g$ and $h_\ell + c = g''$, we thus have $g \sim_{\mathsf{sp}(i)} g''$ and hence $g'' \in \Phi_i(x)$. □

By Lemma 8 we may identify each e-node x in H_i with the corresponding e-node $\hat{x}$ in $\hat{H}'_i$.

We then prove the following lemma.

Lemma 9. *Let x and y be two e-nodes in H_i, and let $\hat{x}$ and $\hat{y}$ be two e-nodes in $\hat{H}'_i$ corresponding to x and y, respectively. Then, $\hat{x}\hat{y} \in E(\hat{H}'_i) = E(H^{c_1}_{i-1}) \cup E(H^{c_2}_{i-1})$ if and only if $xy \in E(H_i)$.*

Proof. We first prove the only-if-part. Suppose that $\hat{x}\hat{y} \in E(H^{c_1}_{i-1})$; it is symmetric for the other case where $\hat{x}\hat{y} \in E(H^{c_2}_{i-1})$. Then, there exist two adjacent nodes $h_x \in \Phi_{i-1}(\hat{x})$ and $h_y \in \Phi_{i-1}(\hat{y})$ such that $h_x(\mathsf{sp}(i-1)) \neq c_1$ and $h_y(\mathsf{sp}(i-1)) \neq c_1$. Therefore, by Lemma 8 we have $h_x + c_1 \in \Phi_i(x)$ and $h_y + c_1 \in \Phi_i(y)$. Note that, since $\hat{x} \neq \hat{y}$, we know that only the spine vertex $\mathsf{sp}(i-1) = \mathsf{sp}(i)$ is recolored between h_x and h_y. Thus, $h_x + c_1$ and $h_y + c_1$ are adjacent in R^0_i, and hence we have $xy \in E(H_i)$.

We then prove the if-part. Since $xy \in E(H_i)$, there exist two adjacent nodes $g_x \in \Phi_i(x)$ and $g_y \in \Phi_i(y)$ in R^0_i. Since $x \neq y$, only the spine vertex $\mathsf{sp}(i)$ is recolored between g_x and g_y, and hence $g_x(v_i) = g_y(v_i)$. Therefore, $g_x[V_{i-1}]$ and $g_y[V_{i-1}]$ are adjacent. We assume that $c_1 = g_x(v_i) = g_y(v_i)$ without loss of generality. Then, since v_i and $\mathsf{sp}(i)$ are adjacent, $g_x(\mathsf{sp}(i)) \neq c_1$ and $g_y(\mathsf{sp}(i)) \neq c_1$. By Lemma 8 we have $g_x \in \bigcup\{\Phi_{i-1}(\hat{x}) \oplus c : c \in L(v_i) \setminus \{\mathsf{col}_{i-1}(\hat{x})\}\}$; this implies that $g_x[V_{i-1}] \in \Phi_{i-1}(\hat{x})$. Similarly, $g_y[V_{i-1}] \in \Phi_{i-1}(\hat{y})$. Since $g_x[V_{i-1}]$ and $g_y[V_{i-1}]$ are adjacent, $\hat{x}\hat{y} \in E(H_{i-1})$. Furthermore, since $\mathsf{col}_{i-1}(\hat{x}) \neq c_1$ and $\mathsf{col}_{i-1}(\hat{y}) \neq c_1$, we have $\hat{x}\hat{y} \in E(H^{c_1}_{i-1})$. Therefore, $\hat{x}\hat{y} \in E(H^{c_1}_{i-1}) \subseteq E(\hat{H}'_i)$. □

We now prove the following lemma.

Lemma 10. $\hat{H}_i = H_i$.

Proof. Recall that H_i consists of a single connected component which contains the e-node z such that $f_0[V_i] \in \Phi_i(z)$. Consider the set of all e-nodes $\hat{x}$ in $\hat{H}'_i$ that correspond to the e-nodes x in H_i. By Lemma 9 the e-node set forms a connected subgraph of a single connected component in $\hat{H}'_i$. Furthermore, by Lemma 8 the component in $\hat{H}'_i$ contains the e-node $\hat{z}$ such that $f_0[V_i] \in \Phi_i(\hat{z})$; by the construction, $\hat{z}$ is contained in $\hat{H}_i$. We thus have $V(H_i) \subseteq V(\hat{H}_i)$.

Therefore, to show $H_i = \hat{H}_i$, by Lemma 9 it suffices to prove that there exists no edge $x\hat{y} \in E(\hat{H}_i)$ which joins two e-nodes $x \in V(H_i)$ and $\hat{y} \in V(\hat{H}_i) \setminus V(H_i)$. Suppose for a contradiction that there exists such an edge $x\hat{y} \in E(\hat{H}_i)$. By the construction, $x\hat{y} \in E(H^{c_1}_{i-1}) \cup E(H^{c_2}_{i-1})$; we may assume that $x\hat{y} \in E(H^{c_1}_{i-1})$ without loss of generality. Then, there exist two adjacent nodes $h_x \in \Phi_{i-1}(x)$ and $h_{\hat{y}} \in \Phi_{i-1}(\hat{y})$ such that $h_x(\mathsf{sp}(i-1)) \neq c_1$ and $h_{\hat{y}}(\mathsf{sp}(i-1)) \neq c_1$. Therefore h_x and $h_{\hat{y}}$ can be extended by c_1, and $h_x + c_1 \in \Phi_i(x)$ and $h_{\hat{y}} + c_1 \in \Phi_i(\hat{y})$ are adjacent in R^0_i. By Lemma 8 we then have $\hat{y} \in V(H_i)$; this contradicts the assumption that $\hat{y} \in V(\hat{H}_i) \setminus V(H_i)$. □

Finally, we show the following lemma.

Lemma 11. $\hat{\mathsf{col}}_i = \mathsf{col}_i$, $\hat{\mathsf{ini}}_i = \mathsf{ini}_i$ *and* $\hat{\mathsf{tar}}_i = \mathsf{tar}_i$.

Proof. Recall that $\mathsf{sp}(i) = \mathsf{sp}(i-1)$ if $v_i \in V_{\mathrm{L}}$. Then, the lemma follows immediately from Lemma 8. □

This completes the proof of Lemma 5.

A.2 Proof of Lemma 6

In this subsection, we prove Lemma 6. Therefore, suppose that we have already computed $(H_{i-1}, \mathsf{col}_{i-1}, \mathsf{ini}_{i-1}, \mathsf{tar}_{i-1})$ for $i \geq 2$, and assume that $v_i \in V_{\mathrm{S}}$.

We first prove that $V(H_i) \subseteq V(\hat{H}'_i)$ holds, as in the following lemma.

Lemma 12. *Let x be any e-node in H_i with $\mathsf{col}_i(x) = c$. Then, there exists exactly one connected component H in H^c_{i-1} such that $\bigcup\{\Phi_{i-1}(y) \oplus c : y \in V(H)\} = \Phi_i(x)$.*

Proof. Let g be any node in $\Phi_i(x) \subseteq V(\mathsf{R}^0_i)$. Then, by Lemma 3 the node $g[V_{i-1}]$ is contained in R^0_{i-1}, and hence there exists an e-node z in H_{i-1} such that $g[V_{i-1}] \in \Phi_{i-1}(z)$. Since $\mathsf{sp}(i) \neq \mathsf{sp}(i-1)$ and $\mathsf{sp}(i)$ is adjacent with $\mathsf{sp}(i-1)$, the assumption $\mathsf{col}_i(x) = c$ implies that $g(\mathsf{sp}(i)) = c$ and hence $g(\mathsf{sp}(i-1)) \neq c$. Therefore, we have $\mathsf{col}_{i-1}(z) \neq c$, and hence H^c_{i-1} has exactly one connected component H that contains z. We thus prove that $\bigcup\{\Phi_{i-1}(y) \oplus c : y \in V(H)\} = \Phi_i(x)$ for the connected component H.

Let g' be any node in $\Phi_i(x)$. Then, $g'(\mathsf{sp}(i)) = c$ and hence it suffices to show that H contains an e-node y such that $g'[V_{i-1}] \in \Phi_{i-1}(y)$. Since $g \sim_{\mathsf{sp}(i)} g'$, there exists a path $\langle g_1, g_2, \ldots, g_\ell \rangle$ in R^0_i such that $g_1 = g$, $g_\ell = g'$ and $g_j(\mathsf{sp}(i)) = g(\mathsf{sp}(i)) = g'(\mathsf{sp}(i)) = c$ for all $j \in \{1, 2, \ldots, \ell\}$. Since $\mathsf{sp}(i)$ is adjacent with $\mathsf{sp}(i-1)$, $g_j(\mathsf{sp}(i-1)) \neq c$ holds for all $j \in \{1, 2, \ldots, \ell\}$. Therefore, there exists a connected component H' in H^c_{i-1} such that the path $\langle g_1[V_{i-1}], g_2[V_{i-1}], \ldots, g_\ell[V_{i-1}] \rangle$ is contained in $\bigcup\{\Phi_{i-1}(y') : y' \in V(H')\}$. Because $z \in V(H)$ and $g_1[V_{i-1}] = g[V_{i-1}] \in \Phi_{i-1}(z)$, we have $H' = H$. Thus, H contains an e-node y such that $g'[V_{i-1}] \in \Phi_{i-1}(y)$. Then, we have $g' \in \bigcup\{\Phi_{i-1}(y) \oplus c : y \in V(H)\}$.

Conversely, let g'' be any node in $\bigcup\{\Phi_{i-1}(y) \oplus c : y \in V(H)\}$. Since H is connected, the subgraph of R^0_{i-1} induced by $\bigcup\{\Phi_{i-1}(y) : y \in V(H)\}$ is connected, too. Then, the induced subgraph contains a path $\langle h_1, h_2, \ldots, h_\ell \rangle$ such that $h_1 = g[V_{i-1}]$ and $h_\ell = g''[V_{i-1}]$. Furthermore, since H is a connected component in H^c_{i-1}, we know that $h_j(\mathsf{sp}(i-1)) \neq c$ for all nodes h_j, $j \in \{1, 2, \ldots, \ell\}$. Therefore, we can extend each node h_j by c, and obtain a path $\langle h_1 + c, h_2 + c, \ldots, h_\ell + c \rangle$. Since $h_1 + c = g$ and $h_\ell + c = g''$, we thus have $g \sim_{\mathsf{sp}(i)} g''$. Since $g \in \Phi_i(x)$, we have $g'' \in \Phi_i(x)$. □

For each e-node $x \in V(H_i)$, let H_x be the connected component in $H^{\mathsf{col}_i(x)}_{i-1}$ which satisfies Lemma 12. Then, we can identify the e-node x in H_i with the e-node $\hat{x}$ in $\hat{H}'_i$ such that $\mathsf{EN}(\hat{x}) = V(H_x)$.

We then prove the following lemma.

Lemma 13. *Let x and y be two e-nodes in H_i, and let $\hat{x}$ and $\hat{y}$ be two e-nodes in $\hat{H}'_i$ corresponding to x and y, respectively. Then, $\mathsf{EN}(\hat{x}) \cap \mathsf{EN}(\hat{y}) \neq \emptyset$ if and only if $xy \in E(H_i)$.*

Proof. We first prove the only-if-part. Suppose that the set $\mathsf{EN}(\hat{x}) \cap \mathsf{EN}(\hat{y})$ contains an e-node a in H_{i-1}. Choose an arbitrary node $h \in \Phi_{i-1}(a)$, then two nodes $h + \mathsf{col}_i(x) \in \Phi_i(x)$ and $h + \mathsf{col}_i(y) \in \Phi_i(y)$ are adjacent in R_i^0. Thus, $xy \in E(H_i)$.

We then prove the if-part. Suppose that $xy \in E(H_i)$, then there exist two adjacent nodes $g_x \in \Phi_i(x)$ and $g_y \in \Phi_i(y)$ in R_i^0. Since g_x and g_y are adjacent and $\mathsf{col}_i(x) \neq \mathsf{col}_i(y)$, we know that $g_x[V_{i-1}] = g_y[V_{i-1}]$. By Lemma 12, there exists an e-node $a \in \mathsf{EN}(\hat{x})$ such that $g_x[V_{i-1}] \in \Phi_{i-1}(a)$. Similarly, there exists an e-node $b \in \mathsf{EN}(\hat{y})$ such that $g_y[V_{i-1}] \in \Phi_{i-1}(b)$. Since $g_x[V_{i-1}] = g_y[V_{i-1}]$, we have $a = b$. Therefore, $a = b \in \mathsf{EN}(\hat{x}) \cap \mathsf{EN}(\hat{y}) \neq \emptyset$. □

By Lemmas 12 and 13, we have $H_i \subseteq \hat{H}'_i$ and $\mathsf{col}_i = \hat{\mathsf{col}}_i$.

We now prove the following lemma.

Lemma 14. *$\mathsf{ini}_i = \hat{\mathsf{ini}}_i$ and $\mathsf{tar}_i = \hat{\mathsf{tar}}_i$.*

Proof. We prove only $\mathsf{ini}_i = \hat{\mathsf{ini}}_i$; it is similar to prove $\mathsf{tar}_i = \hat{\mathsf{tar}}_i$.

Let x be any e-node in H_i such that $\mathsf{ini}_i(x) = 1$. Then, $f_0[V_i] \in \Phi_i(x)$, and $\mathsf{col}_i(x) = \hat{\mathsf{col}}_i(x) = f_0(v_i)$ holds. By Lemma 12 there exists an e-node $y \in \mathsf{EN}(x)$ such that $f_0[V_{i-1}] \in \Phi_{i-1}(y)$ and $\mathsf{ini}_{i-1}(y) = 1$. Since $f_0[V_i] = f_0[V_{i-1}] + \mathsf{col}_i(x)$, we thus have $\hat{\mathsf{ini}}_i(x) = 1$.

Conversely, let $\hat{x}$ be any e-node in $\hat{H}'_i$ such that $\hat{\mathsf{ini}}_i(\hat{x}) = 1$. Then, $\mathsf{EN}(\hat{x})$ contains an e-node y such that $f_0[V_{i-1}] \in \Phi_{i-1}(y)$ and $\mathsf{ini}_{i-1}(y) = 1$. Note that y is in $H_{i-1}^{\mathsf{col}_i(\hat{x})}$, and $f_0[V_{i-1}] + \mathsf{col}_i(\hat{x}) = f_0[V_i]$. By Lemma 12 we thus have $f_0[V_i] \in \Phi_i(\hat{x})$ and hence $\mathsf{ini}_i(\hat{x}) = 1$. □

Finally, we show following lemma.

Lemma 15. *$\hat{H}_i = H_i$.*

Proof. Recall that H_i consists of a single connected component which contains the e-node z such that $f_0[V_i] \in \Phi_i(z)$. By Lemma 13 H_i is contained in one connected component of $\hat{H}'_i$ as a subgraph, and the connected component contains an e-node $\hat{z}$ such that $f_0[V_i] \in \Phi_i(\hat{z})$. By Lemma 14 we have $\hat{\mathsf{ini}}_i(\hat{z}) = 1$, and hence the connected component is indeed $\hat{H}_i$. We thus have $H_i \subseteq \hat{H}_i$.

Therefore, to show $H_i = \hat{H}_i$, by Lemma 13 it suffices to prove that there exists no edge $x\hat{y} \in E(\hat{H}_i)$ which joins two e-nodes $x \in V(H_i)$ and $\hat{y} \in V(\hat{H}_i) \setminus V(H_i)$. Suppose for a contradiction that there exists such an edge $x\hat{y} \in E(\hat{H}_i)$. Then, the set $\mathsf{EN}(x) \cap \mathsf{EN}(\hat{y})$ contains an e-node z. Choose an arbitrary node $h \in \Phi_{i-1}(z)$, then two nodes $h + \mathsf{col}_i(x)$ and $h + \mathsf{col}_i(\hat{y})$ are adjacent in R_i^0. Then, $h + \mathsf{col}_i(\hat{y}) \in \bigcup_{z' \in \mathsf{EN}(\hat{y})} (\Phi_{i-1}(z') \oplus \mathsf{col}_i(\hat{y}))$. By Lemma 12 the corresponding e-node y should be contained in H_i; this contradicts the assumption that $\hat{y} \in V(\hat{H}_i) \setminus V(H_i)$. □

This completes the proof of Lemma 6.

A.3 Proof of Lemma 7

In this subsection, we prove Lemma 7. Since $|V(\hat{H}_i)| \leq |V(\hat{H}'_i)|$ holds, it suffices to prove the following inequality: for each index $i \in \{1, 2, \ldots, n\}$,

$$|V(\hat{H}'_i)| \leq \begin{cases} 2 & \text{if } i = 1; \\ |V(H_{i-1})| + d(v_i) & \text{otherwise.} \end{cases} \tag{1}$$

Consider the case where v_i is a leaf. Then, $V(\hat{H}'_i) = V(H_{i-1})$, and hence Eq. (1) clearly holds. In the remainder of this subsection, we thus consider the case where v_i is a spine vertex.

For a graph $G = (V, E)$, we denote by $\mathsf{cc}(G)$ the number of connected components in G. For a connected graph G, that is, $\mathsf{cc}(G) = 1$, we denote by T_G any spanning tree of G. Since $E(T_G) \subseteq E(G)$, we clearly have the following proposition.

Proposition 1. *Let G be a connected graph, and let V_0 be any vertex subset of G. Then, $\mathsf{cc}(G[V_0]) \leq \mathsf{cc}(T_G[V_0])$ holds.*

We now apply Case (B) of our algorithm to any spanning tree $T_{H_{i-1}}$ of H_{i-1}, instead of applying the operation to H_{i-1}. Let $\hat{H}_i^T$ be the obtained encoding graph, instead of $\hat{H}'_i$. Then, we have the following lemma.

Lemma 16. $|V(\hat{H}'_i)| \leq |V(\hat{H}_i^T)|$.

Proof. For each color $c \in L(v_i)$, let $H_{i-1}^{T,c}$ be the subgraph of $T_{H_{i-1}}$ obtained by deleting all e-nodes y in $T_{H_{i-1}}$ with $\mathsf{col}_{i-1}(y) = c$. Then,

$$|V(\hat{H}'_i)| = \sum_{c \in L(v_i)} \mathsf{cc}(H_{i-1}^c),$$

and

$$|V(\hat{H}_i^T)| = \sum_{c \in L(v_i)} \mathsf{cc}(H_{i-1}^{T,c}). \tag{2}$$

By Proposition 1 we have $\mathsf{cc}(H_{i-1}^c) \leq \mathsf{cc}(H_{i-1}^{T,c})$ for each color $c \in L(v_i)$, and hence $|V(\hat{H}'_i)| \leq |V(\hat{H}_i^T)|$. □

We finally show the following lemma, which verifies Eq. (1) and hence completes the proof of Lemma 7.

Lemma 17. *If v_i is a spine vertex, then $|V(\hat{H}'_i)| \leq |V(H_{i-1})| + d(v_i)$.*

Proof. We first consider the case where $|V(H_{i-1})| = 1$. Then, $\mathsf{cc}(H_{i-1}^c) \leq 1$ for any color $c \in L(v_i)$, and hence

$$|V(\hat{H}'_i)| = \sum_{c \in L(v_i)} \mathsf{cc}(H_{i-1}^c) \leq |L(v_i)|.$$

By Lemma 2 we have $|L(v_i)| \leq d(v_i) + 1$, and hence

$$|V(\hat{H}'_i)| \leq 1 + d(v_i) = |V(H_{i-1})| + d(v_i).$$

We then consider the case where $|V(H_{i-1})| \geq 2$. Recall that $T_{H_{i-1}}$ is a spanning tree of H_{i-1}, and hence $V(T_{H_{i-1}}) = V(H_{i-1})$. For each color $c \in L(v_i)$, let $X_{i-1}(c) = \{x \in V(T_{H_{i-1}}) : \mathsf{col}_{i-1}(x) = c\}$. For each vertex $x \in V(T_{H_{i-1}})$, we denote by $d(T_{H_{i-1}}, x)$ the degree of x in $T_{H_{i-1}}$. Then, by deleting x from $T_{H_{i-1}}$, the number of connected components in the resulting graph is increased by $d(T_{H_{i-1}}, x) - 1$. We thus have

$$\begin{aligned}\mathsf{cc}(H^{T,c}_{i-1}) &= \mathsf{cc}(T_{H_{i-1}}) + \sum\Big\{d(T_{H_{i-1}}, x) - 1 : x \in X_{i-1}(c)\Big\} \\ &= 1 + \sum\Big\{d(T_{H_{i-1}}, x) - 1 : x \in X_{i-1}(c)\Big\}. \end{aligned} \tag{3}$$

By Lemma 16 and Eq. (2) we have

$$|V(\hat{H}'_i)| \leq |V(\hat{H}^T_i)| = \sum_{c \in L(v_i)} \mathsf{cc}(H^{T,c}_{i-1}).$$

Therefore, by Eq. (3) we have

$$\begin{aligned}|V(\hat{H}'_i)| &\leq \sum_{c \in L(v_i)} \Big(1 + \sum\Big\{d(T_{H_{i-1}}, x) - 1 : x \in X_{i-1}(c)\Big\}\Big) \\ &\leq |L(v_i)| + \sum\Big\{d(T_{H_{i-1}}, x) - 1 : x \in V(T_{H_{i-1}})\Big\} \\ &= |L(v_i)| + 2|E(T_{H_{i-1}})| - |V(T_{H_{i-1}})|. \end{aligned} \tag{4}$$

Since $T_{H_{i-1}}$ is a tree, $|E(T_{H_{i-1}})| = |V(T_{H_{i-1}})| - 1$. Furthermore, recall that $T_{H_{i-1}}$ is a spanning tree of H_{i-1}, and hence $V(T_{H_{i-1}}) = V(H_{i-1})$. By Eq. (4) we thus have $|V(\hat{H}'_i)| \leq |L(v_i)| + |V(H_{i-1})| - 2$. By Lemma 2 we have $|L(v_i)| \leq d(v_i) + 1$, and hence

$$|V(\hat{H}'_i)| \leq |L(v_i)| + |V(H_{i-1})| - 2 \leq |V(H_{i-1})| + d(v_i) - 1,$$

as required. □